To what extent is embodied AI necessary (or sufficient) in furthering our understanding of complex adaptive behaviours?

Simulation of Adaptive Behaviour
Alice Eldridge
ace21@cogs

April 22, 2002

Abstract

There is an increasing predilection, amongst all schools of the cognitive sciences, for the consideration of the role of embodiment in complex adaptive behaviours. The role of research in the Simluation of Adaptive Behaviour (SAB), in establishing the importance of embodiment in cognition per say is first addressed. A review of research using real robotics and simulated agents follows, with consideration of the way on which this 'embodied approach' reframes the mainstream AI explainations of cognition. It is concluded that although the field is making promising advances, there are still gaps in the explanatory powers of the new approach.

1 Introduction

Consideration of the importance of embodiment has become increasingly frequent not only within adaptive systems research, but philosophy, psychology, neuroscience, education cognitive anthropology, linguistics and dynamical systems approaches to behaviour and thought. The importance of environment-agent interactions was central in the field cybernetics [1] (arguably the predecssor to contemporary adaptive systems research), and is by no means absent in philosophical considerations of mind [29], [16]. However for several decades, persuaded by the promise of the computational metaphors promoted by GOFAI, the cognitive sciences have operated under a principle of explanatory disembodiment: " the details of the physical embodiment of the agent can safely be ignored for the purposes of cognitive explanation" [42]. Currently, an increasing number of practitioners believe that "cognitive science can on longer afford the simplifications that take the real world and acting agent out of the loop" [8] chapter 3. The congruence in current conception of mind in cognitive science and SAB raises the prospects of the animat researcher achieving its long term goal of "contributing to our understanding of human cognition" [12] p.3.

In many cases examination of the interactions between 'mind' 'body' and 'brain' demand a reformulation of previously central 'mentalistic' notions such as 'representation' and 'concept'. The focus on the importance on 'embeddedness' is indicative of a more general shift from sequential delimited processes, to a more integrated approach. Consideration of functional modules carrying out the 'sense-model-plan-act' cycle of mainstream AI have been replaced by discussion of continuous action-perception and sensorimotor cycles of activity: the attendant Cartesianism is discarded.

Despite the renewed philosophical and theoretical fervor, empirical treatment of the role of embodiment in cognition is difficult and scarce. In this paper the contribution of SAB in establishing the importance of embodiment in cognition, is considered. An illustrated examination of the potentials and problems of both real world robots and simulated embodied agents follows as a means of considering whether cognitive research should adopt the embodied approach exclusively.[1]

2 Establishing the Role of Embodiment in Cognition

According to an overall bottom-up strategy, the fundamental issue is to investigate how the high level cognitive abilities of humans depend upon the evolution of the simplest adaptive abilities and behaviours of animals. In assuming an embodied approach then, the implicit assumption is that high-level cognition is not fundamentally different from the lower-level processing involved in spatial motor, and social interactions. The supposition follows that understanding can be advanced by investigating how the spatial/ motor mechanisms can be used to construct/ evolve and understand 'higher' cognitive capabilities.[25] . In order to validate the approach, this assumption must be verified. Two possible routes of investigation are the analysis of existing species and/ or developmental stages and /or pathologies, in search of key (functional) differences between low and high level cognition or the synthesis of different alternatives to cognition, embodied and otherwise and the study their comparative properties.

Neuroscientific research demonstrates the extent to which the periphery imposes opportunities and constraints for the nervous system, and promotes postulation of shared mechanisms for low- level control of embodied action (eg motor plans for limbs) and higher-level cognition (eg abstract plans). Properties of the periphery simplify neural control, for example the stiffness of muscle around a joint creates an equilibrium point to which a limb returns after perturbation which can be exploited by the spinal cord [5]. The importance of feedback from the body and environment in neural development has also been demonstrated [40].

However, it is difficult to address role of embodiment in more complex adaptive behaviours in natural species, as those species in which we can endorse 'diminishment of bodily capacities', are not typically considered to be capable of higher level cognition. Conversely, we cannot perform structural ablation on higher primates, and human data on principally cognitive deficits is largely anecdotal and sparse. Although promising advances have been made in studying mind-body connection [19], [11], there is currently insufficient data to address the role of embodiment in cognition.

The one empirical investigative avenue for comparative studies of cognitive deficits and capacities is human developmental research: the gradual development and acquisition of both motor and cognitive skills in children provides a healthy corpus of data. However, biological data is largely incomplete, due to inability to experiment with arbitrary systems, and where data exists, it describes disparate cases from which we cannot currently draw conclusions.

The synthetic approach overcomes both of these problems, as there is free reign over what capabilities are included, and what biological properties are modelled. Until we are better able to analyse the biological evidence, a synthetic approach is the best first step in attempting to answer at least the question of how if it is possible that higher-level reasoning can depend upon or be a simple extension of lower-level structures, and thus verify the role of embodiment in cognition per say.

2.1 Embodied Agents

Through their legacy of designer-imposed modularity which separated low and high level reasoning, previous GOFAI approaches implicitly answered this question in the negative. Behaviour Based Robotics (BBR) [7], presents an alternative, providing a means of addressing the question directly. Mataric and colleagues have investigated which aspects of embod-

[1] The entire project demands that we trust that the simulation' of cognitive process , in absence of experience, is profitable. The scientific application of embodied agents is of principle concern, rather than issues surrounding the debate over 'strong AI'.

iment afford and facilitate specific behaviours, and how those behaviours in turn, enable constrain and facilitate higher cognition, by utilising a biologically-inspired notion of basis behaviours as a substrate for control and learning. For example in a navigation task, homing behaviour achieves the goal of getting the robot to the home location, and a wall following behaviour maintains the goal of staying within a range of distances from a wall [25]. The behaviours are directly tied to the dynamics of interaction with the world: they are designed so as to rely on and exploit those dynamics. The specifics of the wall-following behaviour may vary greatly across different robots and environments. In this sense embodiment is crucial; the detailed specifics of the behaviour cannot be derived, tested, or verified outside of its environment of application.

The simulation provides a concrete example of Hendriks-Jansen's notion of *interactive emergence* [15], which is closely tied to ethological concept of *Umwelt* [2]. The low level wall-following behaviour is the product of the interaction between simple reflexes and the environment, but also exists as a new structure in the agents *Umwelt*. The higher level navigational behaviour is entirely dependent upon, and can only exist in, a world 'brought forth' by the lower level activities. The simulation demonstrates not only the importance of embodiment in enabling environmental interactions that facilitate 'cognitive' processes, but illustrates the way in which the approach promotes a 'bottom-up' conceptualisation of complex behaviour.

The role of embodiment in social behaviour has also been demonstrated. [27] Work with embodied agents has shown that the behaviours that served as stable and effective basis for multi-agent interactions were firmly grounded in the immediate, local, physical interactions between agents. This was true of the 'basis set' behaviours including safe-wandering, aggregating, dispersing, following and homing, as well as comparatively higher-level social behaviours, yielding and communicating. It is perhaps surprising that in learning to communicate - to share information about the world - the agents needed repeated physical interactions with each other and with the objects in order for the difficult credit assignment problem to be resolved in the multi-agent domain.

[2]Loosely speaking, the world in which the creature lives.

2.2 Disembodied Multi-agent Systems

Mataric and colleagues are also pursing work with disembodied multi-agent systems consisting of abstract computational agents. For example Kitts [21] developed a 4,096-agent simulation to study resource sharing between heterogeneous spatially distributed agents to determine under what conditions cooperation would emerge and when it would be the most effective/ stable strategy. More abstract information agents have also been explored. The embodied and disembodied multi-agent efforts are being pursued in parallel in search of common principles underlying group dynamics. Comparison between the two domains allows investigation of the role of embodiment in cognition. [25]

3 Real World Robots

Physical realisation can facilitate solutions that are impenetrable when considered from an information processing perspective. Scheier and Lambrinos [36] armed a robot with a gripper and CCD camera, enabling an active vision system. A successful categorisation mechanism was demonstrated, based on a learned re-entrant (ie reciprocally connected) mapping between visual and haptic feature maps, "thus categories develop from the real-time correlations that exist across the independent stimuli. "p2 ibid. These results provide converging evidence in support of child developmental theories which posit that infants discover object categories through the cross-correlation of multimodal experiences [38]. Such work demands a reconceptualisation of categorisation as a sensory-motor coordination rather than a disjunct cognitive subsystem.

As well as the grounding of categorisation in lower level abilities, the role of embodiment in reducing the complexity of a problem has been illustrated. Scheier and Pfeifer [35] describe an autonomous mobile agent that recodes higher dimensional data input by exploiting its physical form and interactions with environment. Correlated data, which can be easily learned due to redundancies resulting from appropriate interactions with environment, is generated. Through this interaction, a type-1 problem is transformed into a type-2 problem, thereby reducing complexity of learning task. In contrast to traditional methods of dealing with increases in difficulty of learning tasks (ie improving the learning algorithm), the use of a real world agent demonstrates that 'cognitive' capacities can only be properly considered in the context of the agents' body and its' environment.

The necessity of physical embodiment is forcefully decried by Brooks : "an agent existing only in simulation would not be complete" [7]. From this stance, physical embodiment is necessary in the study of cognition, as cognition is seen to be embodied in the control architecture of the sensing and acting machine. This is similar to the idea of 'mechanistic embodiment' described by Sharkey and Ziemke [32]and related to the work of Sherrington on reflexes [33] and Loeb on tropisms [23].

Perhaps the earliest significant instantiation of this approach is the work of Grey Walter [13], [14]. Walter's electromechanical tortoises were equipped with two 'sense reflexes' (a small artificial nervous system built from electrical components) which were operated by two 'receptors': one photoelectric, subserving phototactic behaviour, and an electrical contact that served as a touch receptor. The artificial tortoises were attracted to light of moderate intensity, repelled by obstacles, bright light and steep gradients, and never stood still except when re-charging their batteries. Walter claimed that the tortoises exhibited behaviour such as hunger, goal seeking, self-recognition and mutual recognition. He also carried out the first artificial research on classical conditioning with his CORA system. [14]

The study of such simple systems exhibiting complex behaviours cuts both ways. It is arguably useful in that it prevents over theorising, illustrating how simple such apparently complex behaviours may be. The danger exists, however, that very simple phenomenon may be mistaken for something much more complex. The approach encourages a terrible predilection amongst proponents of embodied research to talk of 'emergence of intelligent behaviour', without getting any closer to an understanding of the mechanistic substrate of the comparable behaviour in the living organism. For example, imagine the case of three khepera, each wired to approach light with velocity proportional to light intensity. Inevitably, the situation will arise that two of them join, one crowding another to a side. At this point, the third will sprint over, knocking them apart, 'rescuing' the bullied from the bully. What exactly 'emerges'? Not the rescue attempt, as this exists only in the eye of the beholder. All that really happens is the 'sprint', resulting from the simple mechanism that regulates motor speed as a function of light intensity, coupled with the fact that two lights are brighter than one.

This example highlights the need for caution over the interpretation of and assumed explanatory power of such experiments: Grey Walter Tortoises only exhibited 'goal-seeking ', because lights were placed in places appropriate to elicit the desired anthropomorphic response. Indeed all BBR scenarios are designed so that movement in interaction with well-placed stimuli resembles - *to the human observer* - the behaviour of an organism. [3]. The point is that the robot's behaviour has meaning only for the observer, and there is a danger, when that observer is the expectation-filled designer/experimenter, that the wrong causal significance is attributed to the observed behaviour.

It could be argued that this is a caution that is applicable to conception of these models as examples of strong AI, and need not impede their use as scientific models to advance our understanding of complex behaviours. Whilst the approach does force a reformulation over the *problem* of cognition, the danger

[3] Even within evolutionary approaches, the control is simply abstracted one level - in the design of the fitness function

is in interpreting this reformulation of the problem as a new *solution*.

Real-world robots have been applied usefully as physical tests of the plausibility of mechanistic hypotheses about animal behaviour [39], [22]. Webb demonstrated how a mobile (wheeled) robot could selectively localise the sound of male cricket's stridulations by using comparative auditory intensity to directly control motor output. More than being a specifically useful piece of biological modeling, the work arguably provides a general demonstration that component mechanisms (phonotaxis) of complex behaviour such as mate selection can be mechanised through direct auditory-motor connections. No intermediate recognition or decision processes are necessary. In accordance with the underlying thesis, mentalistic concepts are obviated, and complex mechanisms are described through bodily-environment interactions. But do the mechanisms of the robot phonotaxis and 'mate-selection' constitute an understanding of the comparable behaviour in the real cricket? Perhaps a better understanding of sound localisation is achieved, but the response is comparable only at a distal level. Proximally, phototaxis is achieved in the cricket robot by powering its motors with the output from a device comparing auditory intensity. Our understanding of the mechanisms underlying the behaviour at the proximal level is not advanced.

A concern over the fundamental disparity between robots and living organisms has been expressed by Keijzer [17]. He believes that we are premature in focusing on complex behaviours, pointing out that although models may exhibit similar behavioural patterns, simplicity of sensory-motor interactions means that the underlying mechanisms probably differ. If this concern is correct, then robots do not undergo same process of self-organisation as living systems: they constitute allopoetic rather than autopoetic systems [28]. These arguments certainly apply to discussion of robots as 'strong a-life'. In adopting simulations and robotics as a scientific tool, however, simplifications at all levels are necessary. However, there does seem to be discrepancy between, on the one hand recognising the role of physical form in complex adaptive behaviour, but then ignoring the structural form of the body. Arguably, the muscles underlying the limbs are of equal importance, as are the the nerves controlling the muscles, even perhaps the organic makeup of the calcium channels controlling synaptic communication.... There is undeniably a tension between the employment of paraplegic khepera in adaptive systems research, and the central thesis that higher-level behaviours are subserved by low-level mechanisms. Pragmatically simplification at some level is necessary. An important question that must be addressed is at which level this simplification can justifiably be made.

If on one hand BBR has been charged with over complexity at the proximal level, many believe that the contribution of the approach in advancing our understanding of complex behaviour is seriously impeded by an inability to scale complexity of behaviours at the distal level. Certainly it seems that there will a be a limit to the complexity of hand-designed modules of Brook's approach. An immediate question is, to what extent, if at all, can the embodied approach contribute to our understanding of so-called 'representation-hungry' problem solving [9]. The examples cited above are representative of current research in their focus on active perception, and low-level sensory-motor couplings. There is a paucity of research examining behavioural coordination in the absence of any external stimuli. GOFAI approaches accommodated this 'off- line reasoning' [10] through the use of stored internal representations, but this is precisely this class of mentalistic concept which the embodied bottom-up approach eschews.

An interesting approach is the treatment off-line reasoning as /emphsimulated sensing [37]. Using a mobile robot TOTO, developed by Mataric [26], which encodes geographic information (ie position of walls, corridors etc) in an 'action oriented' way, information about movement and correlated perceptual input (ultrasonic range detectors) are combined. The stored 'map' therefore consists of robotic motion and sonar readings, and as such is perfectly formatted to act as a direct controller, conflating route finding and motion generation into a single act. TOTO is adept at interacting with the local environment, and can, in a weak sense, track situations for which there is no immediate sensory input: it can return, on command, to a previously encountered location. However, the robot is not able to track any location that it has not previously visited.

By using the TOTO architecture off-line, METATOTO is capable of this behaviour. Whereas TOTO uses sonar to act and navigate, METATOTO uses simulated sonar to explore the virtual world. Thus the robots sensors can be stimulated by a floor plan, inducing a sequence of experiences qualitatively similar to those generated by real action. Once the sensors and motors are restored to the real world, online METATOTO can immediately find its way to a target location it has never actually (but virtually) visited.

This work by Stein provides an interesting example of how an embodied agent can solve 'representation-hungry' problems by using the robotic architecture itself. There are of course idealisations: the simulated motion for example is simplified, omitting the irregularities of the real world situation, but the ensuing difficulties could feasibly be rectified in the first instance by adding noise to the simulated route. The scope of this approach is arguably finite however. It is not clear how the sensorimotor simulation strategy can cope with increasingly abstract 'off-line reasoning', such as that concerning implications, rights responsibilities etc. Similar criticisms have been aimed at the work of Brooks by Kirsch [20], who claims that BBR robotics is severely limited, as the robots have no means of forming concepts.

Another major drawback in using physical robots is the preclusion of developmental or evolutionary perspectives. This is noted even by strong proponents of the approach [30] If the necessity of considering the physical form in understanding cognition draws on neuroscientific evidence concerning the role of the body, then evidence from developmental psychology stresses the importance of examining the ongoing interactions between the CNS, organism and environment in the development of complex behaviours.

4 Simulation of Embodied Agents

Beers' work on 'Minimally Cognitive agents' [4], [34] attempts to address both of these issues. He takes an evolutionary approach, using simulated agents rather than physical robots. The simulated agents' behaviour is grounded in an environment. The evolutionary approach adopted means far fewer a priori assumptions need to be made about the necessary design of internal mechanisms. This provides a 'broader intellectual playing field' on which to explore the issues surrounding representation, through the development and analysis of concrete models. The evolved agents exhibit behaviours describable as object discrimination, pointing, [4] object persistence, discrimination of self from non-self, and prediction and recall of future location of objects, [34]. These 'minimally Cognitive agents' exhibit behavioural coordination in the absence of any external stimuli, but to what respect do they advance our understanding of 'cognitive behaviour' and how is the role of embodiment and situatedness demonstrated ?

4.1 Methods of Analysis

Investigation of the interactions between controller, 'body' and environment is not amenable to traditional reductionist or 'componential' analysis. Beer [2] advocates the use of Dynamical Systems theory (DST) as an analytical tool for both simulated and mechanistically embodied agents. In the above example then, this 'cognitive behaviour' is accounted for in terms of dynamic interrelations between behaviour, mechanism and environment: "(the) subtle interplay between sensory input and internal state is crucial to accurate discrimination" (ibid, p.94).

The geometric state-space analysis demands consideration of the coupled system as a whole, rather than dividing analysis into investigation of the network controller and the environment. /footnote-Orthodox cognitive scientist still hold on to a principle of "independence of consciousness and behaviour" [31], a strong motivation for this dogged belief is the notion that to dismiss it is to fall foul of behaviorism. The DST approach offers an attractive alternative in eradicating the mind/body schism and conceiving the agent, its controller and environment as a single system. The behaviour of the agent can then be described in terms of the dynamics of the sub serving mechanisms, obviating the need for recourse to slippery cognitive terms. However, this method of analysis in itself imposes restrictions in terms of tractability and scalability. As the complexity of the dynamical system increases, our intuitive geometric understanding breaks down. In addition, the mathematics of DST becomes decidedly less tractable as the number of parameters and size of state space increase. For example, in applying DST to investigate the operation of a leg controller in a model of insect locomotion citebeer 1995b, the analysis was restricted to a simple five- neuron system controlling one leg. So whilst Beer's explorations of 'minimally cognitive agents' do seem to raise the cognitive capacity of previous embodied agents, there is an inevitable ceiling imposed by the required method of analysis.

The DST approach is undoubtedly useful, and informative in reframing our the analysis constitutes what at first glance seems to be more of an abstract re-description rather than an explanation.

"In its pure form, dynamical explanations make no reference to the structure of the mechanism whose behaviour it is explaining. It tells us how the values of the parameters of the system evolve over time, not what it is about the way the system itself is constituted that causes those parameters to evolve in the specified fashion. It is concerned to explore the topographical structure of the dynamics of the system, but this is a wholly different structure from the system itself" [41]

Thus critics of the DST claim that it does not constitute a full explanation, nor provide adequate understanding of how the larger scale properties are rooted in the interactions of the parts. Even proponents of the approach recognise this shortcoming: Thelen and Smith point out that their depiction of the changing dynamical landscapes of the infants leaves them " completely uninformed about the more precise mechanisms of changing attractor stability." [38].

One solution, independently proposed by Thelen and Smith [38], Kelso [18], and Beer citebeer95, lies in applying the analysis at a range of levels. Beer proposes that we can only begin to comprehend mechanisms of 'cognitive' behaviour by understanding the detailed dynamics of the system at every level, from the interactions between individual neurons, up to the body and environment. This certainly tarries with the bottom-up approach currently advocated, but is not perhaps adhered to in practice by practitioners with an overzealous interest in the agent-environment interactions sub serving more complex behaviours.

5 Conclusion

A major challenge for the embodied approach lies with the class of 'representation-hungry ' problems, and the phenomena of off-line, abstract and environmentally decoupled reason. Although the work of Mataric and Beer in robots and simulated agents respectively has made some promising advance in this area, success to date is restricted to lower-level processes and reasoning. It is important not to conclude, however that the fact that the most convincing demonstrations of the importance of embodiment are to date found in lower-level sensor-motor processes, means that approach necessarily precludes investigation of higher level 'cognition' . In the human case at least, we seem to find at all levels a mixture of highly 'embodied , embedded' strategies and apparently much more abstract and potentially de-coupled strategies, with the creation and manipulation of external symbolic items often functioning as a kind of bridge between the two. So the perceived ceiling in raising the cognitive capacity of agents within the BBR approach could simply be a case of allowing a more complex environment.

Embodiment in the form of Real-world robots has provided startling demonstrations of the extent to which the physics of the body can be exploited to reduce the complexity of the problem, but suffers from scalability problems and precludes investigation of evolution or learning. Simulations of embodied agents exist in a necessarily simplified universe, but relieve practitioners of the need to define excessive assumptions, and enable exploitation of the effects of evolution and learning on structure and behaviour. However, practitioners would be foolish to restrict themselves to embodied agents exclusively. Firstly because in order to establish the importance of embodiment in cognition, comparative studies are needed. In addition, all approaches are restricted in terms of scalability, BBR by hand design, evolutionary robotics by computational power, and all approaches, to be of any worth, must be analysable, which itself imposes restrictions due to tractability. The animat approach to understanding adaptive behaviour has been successful in addressing many of the short comings of previous top down GOFAI approaches. However, the field has yet to provide succeed in the areas of higher cognition championed by their predecessors. To disregard GOFAI approaches entirely would constitute a baby and bathwater scenario. The embodied approach to understanding cognition is undeniably necessary, but perhaps currently insufficient. In the light of their opposing and complementary successes and failures, the most profitable approach to the study of complex adaptive behaviours currently, should be a synthesis of the two.

References

[1] Ashby, W.R. 1960 *Design for a Brain.* Chapman and Hall. London

[2] Beer, R. 1995a Computational and dynamical languages for autonomous agents. In *Mind as Motion*, ed R.Port and T. van Gelder. MIT Press

[3] Beer, R. A Dynamical systems perspective on agent-environment interaction. *Artificial Intelligence* 71: 173-215

[4] Beer, R, 1996,

[5] Bizzi, E., Mussa-Ivaldi, F. and Gister, S. 1991 Science 253, 287 - 291

[6] Brooks, R.A. 1989 A Robot that walks; Emergent Behaviour from a Carefully Evolved Network, *Neural Computation*

[7] Brooks, R., A. 1991 Intelligence Without Reason, in proceedings of IJCAI-91

[8] Clark, A. 1997 *Being There: Putting Brain, body and World Together Again.* MIR Press, Cambridge MA.

[9] Clark, A. and Toribio, J. 1995 Doing without representing? Synthese 101, 401-431

[10] Clark, A. and Grush, R. 2000 Towards a cognitive robotics. Adaptive Behaviour

[11] Damasio, A.R., 1994, *Descartes' Error*, Grosset/ Putnam, NY). Kelso

[12] Guillot, A. and Meyer, J-A. 2001 The animat contribution to cognitive systems research. *Journal of Cognitive System Research* 2, 157-165.

[13] Grey Walter, W. 1950 An imitation of Life. *Scientific American* 182: 42-54

[14] Grey Walter, W. 1953 *The Living Brain* New York: Norton

[15] Hendriks-Jansen, H. 1994 *Situated Activity, Interactive Emergence, and Human Thought* Ph.D Thesis, COGS, University of Sussex, Brighton, England. Published in *Catching Ourselves in the Act: Situated Activity, Interactive Emergence, and Human Thought* MIT Press, 1996.

[16] Heidegger, M. 1962 *Being and Time.* Oxford: Blackwell.

[17] Keijzer, F.A. 1998 Some armchair worries about wheeled behaviour. In *From Animals to Animats 5* pp 13-21 Cambridge, MA: MIT Press.

[18] Kelso, S., 1995 *Dynamic Patterns.* MIT Press.

[19] Kinsbourne, M. 1995, Awareness of One's own body: An Attentional Theory of Its Nature, Development, and Brain Basis, in J.L. Bermudez, A. Marcel and N. Eilan, eds, *The Body and the Self*, MIT Press, Cambridge, MA

[20] Kirsch, D. 1991 Today the Earwig, tommorrow man. *Artificial Intelligence* 47: 161-184

[21] Kitts, B. 1996 The Evolution of Multi-celled organisation, Technical Report working paper, Volen Center of complex systems, Brandeis University.

[22] Lambrinos, D. et al 1997 An autonomous agnet navigating with a polarised light compass. *Adaptive Behaviour* 6(1): 131-161.

[23] Loeb, J. 1918 *Forced Movements, Tropisms and Animal Conduct* Philadelphia Lippincott

[24] Maes, P. 1989, The Dynamics of Action Selection, in IJCAI-89, pp991-997, Detroit, MI.

[25] Mataric, M 1997 Studying the Role of Embodiment in Cognition

[26] Mataric, Maja, 1992, Integration of representation into goal-driven behaviour- based robots IEEE journal of Robotics Automation , 304-312

[27] Mataric 1996b Learning in Multi-Robot Systems in G. Weiss, and S.Sen, eds, 'adaptation and Learning in Multi-Agent Systems', Vol. Lecture notes in Artificial intelligence vol 1042, Springer-Verlag, pp.152-163

[28] Maturana, H.R. and Varela, F.J. 1987 *Autopoesis and Cognition: The Realisation of the Living.* Dordrecht: D. Reidel

[29] Merleau-Ponty, M. 1962 *Phenomenology of Perception.* London: Routledge.

[30] Pfeifer, R. 1996 Building "Fungus Eaters": Design principles of Autonomous agents. *Proceedings of SAB*

[31] Searle, J. 1980 Minds, brains and programs. *Behavioural Brain Sciences* 3: 417-457.

[32] Sharkey, N., E. and Ziemke, T. 2001 Mechanistic versus phenomenal embodiment: Can robot embodiment lead to strong AI? *Journal of Cognitive Systems Research* 2: 251-262

[33] Sherrington, C.S., 1906 *The integrative action of the nervous system*, New York: C. Schriber's sons.

[34] Slocum,A.C., Downey, D.C, and Beer, R. 2000 Further experiments in the evolution of minimally cognitive behaviour: From perceiving affordances to selective attention. *SAB-6*

[35] Scheier, C. and Pfeifer, R. 1998 Exploiting embodiment for category learning.

[36] Scheier and Lambrinos 1998 Categorisation in a real-world agent using haptic exploration and active perception.

[37] Stein, L. (1994) imagination and situated cognition. *Journal of Experimental Theoretical artificial Intelligence* 6. 393-407

[38] Thelen and Smith 1994 A dynamic System approach to the Development of Cognition and Action, MIT Press".

[39] Webb, B. 1995 Using Robots to model animals: A Cricket Test. *Robotic and Autonomous Systems* 16(2-4): 117-134.

[40] Westerga, J. and Gramsbergen, A. 1993 Early Human Development. 34, 89-100

[41] Van Gelder, T.1991, What might cognition be if not computation ? Journal of philosophy 92, no 7: 345-381)

[42] Wheeler, M. 1995 The Philosophy of situated activity. PhD thesis, COGS university Sussex, Falmer, Brighton, UK